

Improved Social Welfare Bounds for GSP at Equilibrium

Brendan Lucier*

Renato Paes Leme†

Abstract

The Generalized Second Price auction is the primary method by which sponsored search advertisements are sold. We study the performance of this auction under various equilibrium concepts. In particular, we demonstrate that the Bayesian Price of Anarchy is at most $2(1 - 1/e)^{-1} \approx 3.16$, significantly improving upon previously known bounds.

Our techniques are intuitively straightforward and extend in a number of ways. For one, our result extends to a bound on the performance of GSP at coarse correlated equilibria, which captures (for example) a repeated-auction setting in which agents apply regret-minimizing bidding strategies. In addition, our analysis is robust against the presence of byzantine agents who cannot be assumed to participate rationally.

Additionally, we present tight bounds for the social welfare obtained at pure NE for the special case of an auction for 3 slots, and discuss potential methods for extending this analysis to an arbitrary number of slots.

1 Introduction

The sale of advertising space is the primary source of revenue for many providers of online services. This is due, in part, to the fact that providers can tailor advertisements to the preferences of individual users. A search engine, for example, can choose to display ads that synergize well with a query being searched. However, such dynamic provision of content complicates the process of selling ad space to potential advertisers. The now-standard method has advertisers place bids – representing the amount they would be willing to pay per click – which are resolved in an automated auction whenever ads are to be displayed.

By far the most popular bid-resolution method currently in use is the Generalized Second Price auction (GSP), a generalization of the well-known Vickrey auction. In the GSP, there are multiple ad “slots” of varying appeal (i.e. slots at the top of the page are more effective). Advertisers are assigned slots in order of their bids, with the highest bidders receiving the best slots; each advertiser then pays an amount equal to the bid of the next-highest bidder. While simple to understand and use, the GSP has some notable drawbacks: unlike the Vickrey auction it is not truthful, and it does not generally guarantee the most efficient outcome (i.e. the outcome that maximizes social welfare). Nevertheless, the use of GSP has been extremely successful in practice. This begs the question: *are there theoretical properties of the Generalized Second Price auction that would explain its prevalence?*

Here we continue the line of work aimed at answering this question by analyzing the performance of GSP under various models of rational agent behaviour. First, we consider Bayes-Nash equilibria (BNE) of GSP. In this model, the auction is viewed as a partial-information game in which each participant’s value per click is private information drawn independently at random from commonly-known distributions. Such a model is particularly relevant for online ad auctions, since historical data can readily be observed to develop accurate market statistics. A BNE is then a profile of bidding strategies whereby each agent maximizes his expected profit subject to the distribution over the other agents’ values. We study the expected social welfare that GSP attains at any BNE, as a fraction of the optimal social welfare.

*blucier@cs.toronto.edu. Dept. of Computer Science, University of Toronto, Toronto, ON.

†renatoppl@cs.cornell.edu. Dept. of Computer Science, Cornell University, Ithaca, NY.

This metric is commonly known as the Bayesian Price of Anarchy, representing the loss in efficiency due to having outcomes determined at BNE rather than a benevolent optimizer.

The BNE solution concept captures scenarios in which a large market of advertisers settle into a stable pattern of bidding strategies. However, empirical studies show that bidding need not stabilize in some cases. Advertising slot auctions can be repeated millions of times per day, and there are bidding patterns in which agents modify their strategies over time to respond to each others' bids. To address such cases, one must consider GSP in the broader context of a repeated auction. In such settings, we assume that an agent's value per click does not change over time, but declared bids can change each round. A solution concept then describes rational behaviour over many instances of the auction (i.e. a method of responding to the past play of other agents), and the metric of interest is the average social welfare attained by GSP over many rounds.

We consider an equilibrium model suited to long-run bidder behaviour in GSP. Namely, we consider settings in which agents choose their bids so that their *regret* vanishes over time. Roughly speaking, such a model assumes that agents observe the bidding patterns of others and modify their own bids in such a way that their long-term performance approaches that of a single optimal strategy chosen in hindsight. It is well-known from learning theory that such regret minimization is easy to achieve via simple bidding techniques. We bound the Price of Total Anarchy, which is the ratio between the social welfare of the optimal allocation and the average social welfare obtained by GSP when agents minimize regret over a sufficiently long number of rounds.

Results Our main result is a bound on the social welfare obtained at Bayes-Nash equilibrium for the GSP auction. Specifically, we show that the Bayesian Price of Anarchy for GSP is at most $2(1 - 1/e)^{-1} \approx 3.164$. This improves upon the previous best-known bounds of 8 for BNE and 4 for (mixed) NE [9].

Perhaps just as important as the improved bounds, however, is the straightforward and robust nature of our proof. In particular, our results extend to give the same bound for coarse correlated equilibria, which implies that the Price of Total Anarchy is at most $2(1 - 1/e)^{-1}$. Moreover, these results are resilient against the presence of Byzantine agents, in the following sense. Suppose that, in addition to the rational participants in the auction, there is also some set of agents who apply arbitrary strategies. We can view these as irrational participants who do not understand how to bid strategically. Note that, in such a setting, it is not possible for an auction to guarantee a fraction of the social welfare obtainable from the irrational bidders; after all, a bidder with very large value may decide (irrationally) to bid 0 and effectively not participate in the auction. What we *can* show, however, is that the presence of the irrational bidders does not interfere with the auction's ability to approximate the welfare obtainable from the rational bidders. That is, the ratio of the optimal social welfare *of the rational bidders* to the total social welfare obtained at any BNE is at most $2(1 - 1/e)^{-1} \approx 3.164$. This result requires an assumption on the play of the irrational bidders, which is that no player bids more than his true value. We feel that this is a reasonable assumption, as overbidding is a dominated strategy that is easily avoided; we therefore view the irrational bidders as novice or uninformed participants who would avoid dominated strategies, rather than truly adversarial agents.

Our results hold for a standard model of separable click-through rates, where the probability that a user clicks on an advertisement j in slot i is of the form $\alpha_i \gamma_j$. That is, it is a product of two separable components: one for the advertiser, and one for the slot. For ease of exposition, we will focus on the special case that $\gamma_j = 1$ for all j . However, we note that our results extend to the more general case of separable click-through rates.

Related work In recent years there has been a surge of work on algorithmic mechanism design for sponsored search, beginning with Mehta et al. [13, 12]. See the survey of Lahaie et al [7] for an overview of subsequent developments. The GSP model applied in this manuscript is due to Edelman et al [2] and Varian [15].

The work most closely related to ours is that of Paes Leme and Tardos, who also study equilibria of GSP [9]. They give upper bounds on the Price of Anarchy in pure, mixed, and Bayesian strategies; achieving bounds of 1.618, 4, and 8, respectively. Our main result is a simplification and strengthening of their results for the mixed and Bayesian cases, as well as an extension to different but related solution concepts.

Edelman et al [2] and Varian [15] study Envy-free equilibria of GSP (a special case of Nash equilibrium) in the full information setting. They demonstrate that such equilibria exist, and that all such equilibria are socially optimal. Gomes and Sweeney [3] study the Generalized Second Price Auction in the Bayesian context. They show that, unlike the full information case, there may not exist symmetric or socially optimal equilibria in this model, and obtain sufficient conditions on click-through-rates that guarantee the existence of a symmetric and efficient equilibrium. Lahaie [8] also considers the problem of bounding the social welfare obtained at equilibrium, but restricts attention to the special case that click-through-rate α_i decays exponentially along the slots with a factor of $\frac{1}{\delta}$. Lahaie proves a price of anarchy of $\min\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\}$.

Lucier and Borodin [11] study the Bayesian price of anarchy for greedy combinatorial auctions. They show via a type of smoothness argument (see [14]) that a greedy c -approximation algorithm can be turned into a mechanism with Price of Anarchy $c + 1$ - for pure and mixed Nash and for Bayes-Nash equilibria. Lucier [10] considers repeated greedy auctions and studies the design of mechanisms with bounded price of total anarchy and price of sinking. These results do not imply bounds for GSP, since it is not a combinatorial auction (and, in particular, GSP does not provide a bidding language expressive enough to implement their mechanisms). However, the approach taken in our work is similar to the one that drives their results.

The study of regret-minimization goes back to the work of Hannan on repeated two-player games [4]. Kalai and Vempala [6] extend the work of Hannan to online optimization problems, and Kakade et al [5] further extend to settings of approximate regret minimization. Blum et al [1] apply regret-minimization to the study of inefficiency in repeated games, coining the phrase “price of total anarchy” for the worst-case ratio between the optimal objective value and the average objective value when agents minimize regret.

2 Preliminaries

We consider an auction with n advertisers and n slots¹. An *outcome* is an assignment of advertisers to slots. An outcome can be viewed as a permutation π with $\pi(k)$ being the player assigned to slot k . Being assigned to the k -th slot results in α_k clicks, where $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$. Each advertiser i has a private type v_i , representing his or her value per click received. The sequence $\mathbf{v} = (v_1, \dots, v_n)$ is referred to as the *type profile*. We will write \mathbf{v}_{-i} for \mathbf{v} excluding the i th entry, so that $\mathbf{v} = (v_i, \mathbf{v}_{-i})$.

A mechanism for this auction elicits a bid $b_i \in [0, \infty)$ from each agent i , which is interpreted as a type declaration, and returns an assignment as well as a price p_i per click for each agent. If advertiser i is assigned to slot j , his *utility* is $\alpha_j(v_i - p_i)$, which is the number of clicks received times profit per click. The *social welfare* of outcome π is $SW(\pi, \mathbf{v}) = \sum_j \alpha_j v_{\pi(j)}$, the total value of the solution for the participants. The optimal social welfare is $OPT(\mathbf{v}) = \max_{\pi} SW(\pi, \mathbf{v})$.

We focus on a particular mechanism, the Generalized Second Price auction, which works as follows. Given bid profile \mathbf{b} , the auction sets $\pi(k)$ to be the advertiser with the k th highest bid (breaking ties arbitrarily). That is, GSP assigns slots with higher click-through-rate to agents with higher bids. Payments are then set according to $p_i = b_{\pi^{-1}(i)+1}$. That is, the payment of the k th highest bidder is precisely the bid of the next-highest bidder (where we take $b_{n+1} = 0$). We will write $u_i(\mathbf{b})$ for the utility derived by agent i from the GSP when agents bid according to \mathbf{b} .

For the remainder of the paper, we will write $\pi(\mathbf{b}, j)$ to be the player assigned to slot j by GSP when the agents bid according to \mathbf{b} . We will also write $\sigma(\mathbf{b}, i)$ for the slot assigned to bidder i by GSP, again when agents bid according to \mathbf{b} . We write $\pi^i(\mathbf{b}_{-i}, j)$ to be the player that would be assigned to slot j by GSP if agent i did not participate in the auction. We will write $\nu(\mathbf{v})$ for the optimal assignment of slots to bidders for value profile \mathbf{v} , so that $\nu(\mathbf{v}, i)$ is the slot that would be allocated to agent i in the optimal assignment².

¹we handle unequal numbers of slots and advertisers by adding virtual slots with click-through-rate zero or virtual advertisers with zero value per click.

²We note that, since GSP makes the optimal assignment for a given bid declaration, we actually have that $\nu(\mathbf{v}, i)$ and $\sigma(\mathbf{v}, i)$ are identically equal. We define ν mainly for use when emphasizing the distinction between an optimal assignment for a value profile and the assignment that results from a given bid profile.

2.1 Pure and Mixed Nash Equilibrium

A (*pure*) *strategy* for agent i is a function $b_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that maps each private value to a declared bid. A *mixed strategy* maps a private value to a distribution over bids, corresponding to a randomized declaration.

We will make the standard assumption that agents apply strategies that never overbid. That is, we restrict our attention to strategies in which $b_i(v_i)$ assigns probability 0 to all bids larger than v_i , for all i and v_i . This assumption is motivated by the fact that overbidding is a dominated strategy: an agent's expected utility can only increase by replacing a bid larger than v_i with a bid of v_i .

Given a value profile \mathbf{v} , we say that strategy profile \mathbf{b} is a mixed Nash equilibrium if, for all i and all alternative strategies $b'_i(\cdot)$,

$$\mathbb{E}u_i(b_i(v_i), b_{-i}(v_{-i})) \geq \mathbb{E}u_i(b'_i(v_i), b_{-i}(v_{-i})).$$

That is, each agent i maximizes his utility by bidding according to strategy $b'_i(\cdot)$. We say this is a pure Nash equilibrium if, in addition, all strategies are pure. We define the (*mixed*) *Price of Anarchy* to be the worst-case ratio between social welfare in the optimum and expected social welfare in GSP across all valuation profiles and all mixed Nash equilibria:

$$\sup_{\mathbf{v}, \mathbf{b}(\cdot) \text{ NE}} \frac{OPT(\mathbf{v})}{\mathbb{E}_{\mathbf{b}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]}.$$

2.2 Bayesian setting

In a Bayesian setting, we suppose that each agent's type is drawn from a publicly known distribution. That is, $\mathbf{v} \sim \mathbf{F}$ where $\mathbf{F} = F_1 \times F_2 \times \dots \times F_n$. We then say that strategy profile \mathbf{b} is a Bayes-Nash equilibrium for distributions \mathbf{F} if, for all i , all v_i , and all alternative strategies b'_i ,

$$\mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \geq \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(b'_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))]$$

That is, each agent maximizes his expected utility by bidding in accordance with strategy $b_i(\cdot)$, where expectation is taken over the distribution of the other agents' types and any randomness in their strategies. We define the *Bayes-Nash Price of Anarchy* to be the worst-case ratio between social welfare in the optimum and social welfare in GSP across all distributions and all Nash equilibria:

$$\max_{\mathbf{F}, \mathbf{b}(\cdot) \text{ BNE}} \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})]}{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{b}(\mathbf{v})}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})]}.$$

2.3 Repeated Auctions

We now turn to repeated versions of GSP. In this setting, the GSP auction is run $T \geq 1$ times with the same slots and agents. The private value profile \mathbf{v} of the agents does not change between rounds, but the agents are free to change their bids. We write b_i^t for the bid of agent i on round t . We refer to $D = (b^1, \dots, b^T)$ as a *declaration sequence*. We will write $\pi(D)$ for the sequence of permutations generated by GSP on input sequence D . The average social welfare generated by GSP is then $SW(\pi(D), \mathbf{v}) = \frac{1}{T} \sum_t SW(\pi(b^t), \mathbf{v})$.

Declaration sequence $D = (b^1, \dots, b^T)$ *minimizes external regret* for agent i if, for any fixed declaration b_i , $\sum_t u_i(b_i^t, b_{-i}^t) \geq \sum_t u_i(b_i, b_{-i}^t) + o(T)$. That is, as T grows large, the utility of agent i approaches the utility of the optimal fixed strategy in hindsight. The *Price of Total Anarchy* is the worst-case ratio between social welfare in the optimum and the average social welfare obtained by GSP across all declaration sequences that minimize external regret for all agents. That is, the price of total anarchy is

$$\lim_{T \rightarrow \infty} \max_{\mathbf{v}, D} \frac{OPT(\mathbf{v})}{SW(\pi(D), \mathbf{v})}$$

where the maximum is taken over declaration sequences that minimize external regret for all agents.

3 Bayesian Price of Anarchy

In this section we prove the following upper bound on the Bayesian Price of Anarchy for GSP.

Theorem 1 *The Bayesian Price of Anarchy of GSP is at most $2(1 - 1/e)^{-1} \approx 3.164$.*

The proof of Theorem 1 proceeds in two steps. We first show that a structural property of bidding profiles implies a bound on the social welfare obtained by GSP (Lemma 2). We then show that this structural property holds at all BNE of the GSP (Lemma 3).

Lemma 2 *Suppose that $\mathbf{v} \sim \mathbf{F}$ and agents apply strategy profile $\mathbf{b}(\cdot)$. Suppose further that the following is true:*

$$\mathbb{E}_{\mathbf{v}_{-i}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i + \alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}), k)}] \geq \gamma \alpha_k v_i \quad \text{for all slots } k, \text{ players } i, \text{ and values } v_i. \quad (1)$$

Then $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2} \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})]$.

Lemma 3 *At any BNE of GSP, (1) holds with $\gamma = 1 - \frac{1}{e}$.*

Lemma 2 and Lemma 3 immediately imply Theorem 1.

Proof of Lemma 2 : Fix some value profile \mathbf{v} . For notational convenience, let Γ be the induced distribution on bid profiles $\mathbf{b} = \mathbf{b}(\mathbf{v})$ when $\mathbf{v} \sim \mathbf{F}$. Then for any player i , value v_i , and slot k , if we write $b_i = b_i(v_i)$, then we can express (1) as:

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(\mathbf{b}, i)} v_i] + \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_k b_{\pi^i(\mathbf{b}_{-i}, k)}] \geq \gamma \alpha_k v_i.$$

Note that \mathbf{v}_{-i} does not appear in this expression; bids \mathbf{b}_{-i} are taken to be drawn from induced distribution Γ_{-i} . Now, recalling that $\nu(\mathbf{v}, i)$ is the slot assigned to player i in the optimal assignment for values \mathbf{v} , we can take $k = \nu(\mathbf{v}, i)$ in the above inequality. We then have

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(\mathbf{b}, i)} v_i] + \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))}] \geq \gamma \alpha_{\nu(\mathbf{v}, i)} v_i$$

for all \mathbf{v} and all i . Notice that the strategy $b_i(\cdot)$ does not appear in the second term, so we can rewrite as

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(\mathbf{b}, i)} v_i] + \mathbb{E}_{\mathbf{b} \sim \Gamma}[\alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))}] \geq \gamma \alpha_{\nu(\mathbf{v}, i)} v_i.$$

Summing over all i and taking expectation over \mathbf{v} :

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] + \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b} \sim \Gamma}[\alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))}] \right] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\nu(\mathbf{v}, i)} v_i \right]$$

Consider each of the three expectations in the above expression. For the third term, we note

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\nu(\mathbf{v}, i)} v_i \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})].$$

For the first term, linearity of expectation implies

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}}[\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})].$$

For the second expectation, notice that:

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{b} \sim \Gamma} \left[\sum_i \alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))} \right] \leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_i \alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))} \right] = \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_k \alpha_k b_{\pi(b, k)} \right]$$

which is $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{v}), \mathbf{v})]$.

We therefore conclude $2\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{v}), \mathbf{v})] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})]$, completing the proof. \blacksquare

Proof of Lemma 3 : We wish to show that $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i + \alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}), k)}] \geq \gamma \alpha_k v_i$, for all slots k , players i , and values v_i . First, note that if $\alpha_k < \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)}]$ then the result is trivial. So, let's consider $\alpha_k \geq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)}]$. We'll prove that

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i + \alpha_k b_{\pi^i(\mathbf{b}_{-i}(\mathbf{v}_{-i}), k)}] \geq \alpha_k v_i - \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i] \cdot \log \frac{\alpha_k v_i}{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i]}$$

and then dividing everything by $\alpha_k v_i$ and using that $\frac{\log(x)}{x} \leq \frac{1}{e}$ we get the desired result.

Consider any bid b'_i for agent i such that $b_i < b'_i < v_i$. Then for each slot k , since bid b_i is utility-maximizing for agent i , the utility of bidding b'_i satisfies

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i] \geq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[u_i(b'_i, \mathbf{b}_{-i})].$$

Also, if agent i bids b'_i and moreover it is true that $b_{\pi^i(\mathbf{b}_{-i}, k)} < b'_i$, then agent i will win a slot with at least α_k clicks. Thus,

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[u_i(b'_i, \mathbf{b}_{-i})] \geq (v_i - b'_i) \alpha_k \mathbb{P}_{\mathbf{v} \sim \mathbf{F}}[b_{\pi^i(\mathbf{b}_{-i}, k)} < b'_i].$$

Combining these two inequalities and substituting $z = v_i - b'_i$, we get

$$\mathbb{P}_{\mathbf{v} \sim \mathbf{F}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} > z] \leq \frac{1}{z \cdot \alpha_k} \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i].$$

We are now able to estimate the expected value of $v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}$ using the fact that

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}] \geq \int_0^\infty \mathbb{P}_{\mathbf{v} \sim \mathbf{F}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} > z] dz.$$

Since $v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} \leq v_i$ with probability 1, we have

$$\begin{aligned} \mathbb{E}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)}] &= \int_0^{v_i} \mathbb{P}_{\mathbf{v} \sim \mathbf{F}}[v_i - b_{\pi^i(\mathbf{b}_{-i}, k)} > z] dz \\ &\leq \int_0^{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i] / \alpha_k} 1 dz + \int_{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i] / \alpha_k}^{v_i} \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i]}{\alpha_k z} dz \\ &\leq \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i]}{\alpha_k} + \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i]}{\alpha_k} \left(\log v_i - \log \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\alpha_{\sigma(\mathbf{b}(\mathbf{v}), i)} v_i]}{\alpha_k} \right). \end{aligned}$$

Multiplying both sides by α_k and rearranging gives the required inequality. \blacksquare

3.1 Correlated bids and Price of Total Anarchy

Notice that the proof of the previous section applies even in cases where agent bids are coarsely correlated. In such a case, we can consider a common source of randomness \mathcal{R} and each bidding function to be a function $b_i(v_i, r)$, where $r \sim \mathcal{R}$. We call a profile of bidding functions a coarse correlated equilibrium if:

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[u_i(b_i(v_i, r), \mathbf{b}_{-i}(\mathbf{v}_{-i}, r))] \geq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[u_i(b'_i(v_i, r), \mathbf{b}_{-i}(\mathbf{v}_{-i}, r))], \forall i, v_i, r$$

We still suppose $v_i \sim F_i$ where F_i are independent distributions. In this case, F and \mathcal{R} induce a distribution Γ on the bids.

Adapting Lemma 3 to this context is straightforward. Now, to adapt Lemma 2, observe that the only additional requirement is that we must argue that

$$\mathbb{E}_{\mathbf{b} \sim \Gamma}[\alpha_{\sigma(v, i)} b_{\pi^i(\mathbf{b}_{-i}, \sigma(v, i))}] = \mathbb{E}_{\mathbf{b} \sim \Gamma}[\alpha_{\sigma(v, i)} b_{\pi^i(\mathbf{b}_{-i}, \sigma(v, i))}].$$

However, this follows because the marginal of $b \sim \Gamma$ restricted to $-i$ is exactly $b_{-i} \sim \Gamma_{-i}$.

We now note that the result from the previous section implies a bound on the price of total anarchy for GSP. This follows because, whenever bidding sequence $D = (\mathbf{b}^1, \dots, \mathbf{b}^T)$ minimizes regret for all agents, the bidding strategy with shared randomness $b_i(v_i, t) = b_i^t(v_i)$ for $t \in [T]$ is a coarse correlated equilibrium. Lemmas 2 and 3 therefore imply that, for all \mathbf{v} ,

$$\mathbb{E}_{t \in [T]}[SW(\pi(\mathbf{b}^t(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2}(1 - 1/e)OPT(\mathbf{v})$$

which implies that the price of total anarchy is bounded by $2(1 - 1/e)^{-1}$.

3.2 Byzantine Agents

We now consider a setting in which, in addition to the n advertisers who bid rationally, there are m “byzantine” advertisers who may bid irrationally. Write N for the set of rational advertisers, and M for the set of irrational advertisers. Note that we still think of the irrational advertisers as being true players, who still receive value per click. The irrational bidders simply cannot be assumed to play at equilibrium; for example, they may not have experience with the GSP auction, or not know about historical bidding patterns.

Given an outcome π (which is an assignment of these $n + m$ bidders to $n + m$ slots), the definition of social welfare is unchanged: it is $SW(\pi, \mathbf{v}) = \sum_{i \in N \cup M} v_i \alpha_{\pi(i)}$. We define the social welfare of bidders in N to be precisely that: $SW_N(\pi, \mathbf{v}) = \sum_{i \in N} v_i \alpha_{\pi(i)}$. The optimal social welfare for bidders in N is $OPT_N(\mathbf{v}) = \max_{\pi} SW_N(\pi, \mathbf{v})$.

We wish to show that the total social welfare obtained by GSP is a good approximation to $OPT_N(\mathbf{v})$ when the players in N play at equilibrium and the players in M play arbitrarily. That is, the addition of irrational players does not degrade the social welfare guarantees of GSP had they not participated. In order to make this claim, we must impose a restriction on the behaviour of the irrational players: that they do not overbid. In other words, we require that $b_i(v_i) \leq v_i$ for all $i \in M$ and all v_i . We feel this is a natural restriction: overbidding is easily seen to be a dominated strategy (i.e. any strategy that bids higher than v_i is dominated by a strategy that lowers such bids to be at most v_i). Moreover, it is arguable that inexperienced bidders would bid conservatively, and not risk a large payment with no gain.

Given that byzantine agents do not overbid, we note that our BPOA bounds go through in this setting almost without change. In particular, our structural property (1) continues to hold for all agents in N .

Lemma 4 Equation (1) holds with $\gamma = 1 - \frac{1}{e}$ for all for all $i \in N$.

Proof. This proof follows the proof of Lemma 3 without change. Note that in that proof we used only the fact that the bidding strategy of agent i is a best response, so the fact that other agents may not bid at equilibrium does not affect the argument. ■

The corresponding version of Lemma 2 then follows from (1) just as in the setting without byzantine agents.

Lemma 5 If (1) holds for all $i \in N$, then $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \geq \frac{1}{2}\gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT_N(\mathbf{v})]$.

Proof. Precisely as in the proof of Lemma 2, we obtain

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}} [\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] + \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b} \sim \Gamma} [\alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}, \nu(\mathbf{v}, i))}] \right] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} v_i \right]$$

where we note that the summations are over agents in N . Then, as in Lemma 2,

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} v_i \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[OPT(\mathbf{v})].$$

and

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i \in N} \mathbb{E}_{\mathbf{b}_{-i} \sim \Gamma_{-i}} [\alpha_{\sigma(\mathbf{b}, i)} v_i] \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW_N(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})] \leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[SW(\pi(\mathbf{b}(\mathbf{v})), \mathbf{v})].$$

For the second expectation, notice that:

$$\begin{aligned} \mathbb{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{b} \sim \Gamma} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} b_{\pi^i(\mathbf{b}_{-i}, \nu(\mathbf{v}, i))} \right] &\leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_{i \in N} \alpha_{\nu(\mathbf{v}, i)} b_{\pi(\mathbf{b}, \nu(\mathbf{v}, i))} \right] \\ &= \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_{k: \exists i \in N, \nu(\mathbf{v}, i) = k} \alpha_k b_{\pi(\mathbf{b}, k)} \right] \leq \mathbb{E}_{\mathbf{b} \sim \Gamma} \left[\sum_k \alpha_k v_{\pi(\mathbf{b}, k)} \right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW(\pi(\mathbf{v}), \mathbf{v})]. \end{aligned}$$

We therefore conclude $2\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [SW(\pi(\mathbf{v}), \mathbf{v})] \geq \gamma \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [OPT_N(\mathbf{v})]$, completing the proof. \blacksquare

Together, Lemma 4 and Lemma 5 imply that the Bayesian Price of Anarchy of GSP is at most $2(1 - 1/e)^{-1}$ even in the presence of irrational bidders. Following the comments in Section 3.1, we can apply the same argument to obtain a matching bound on the Price of Total Anarchy with irrational bidders.

4 Towards a Tight Pure PoA

In [9], Paes Leme and Tardos give a bound of 1.618 for the Pure Price of Anarchy of GSP. They also prove that for $n = 2$ slots, the correct bound is 1.25. Here we show that for $n = 3$, the correct bound is 1.25913. We conjecture that this is the correct Price of Anarchy for GSP (for any number of slots) and we suggest an approach to prove this result.

Lemma 6 *For $n = 3$ slots, the pure Price of Anarchy of GSP is exactly 1.25913.*

Proof. Fix one permutation π . If there is an i s.t. $\pi(i) = i$ then it is easy to show the Price of Anarchy is bounded by 1.25. This excludes all but two allocations which we analyze below. They are: (i) $\pi = [2, 3, 1]$ and (ii) $\pi = [3, 1, 2]$.

Case (i): $\pi = [2, 3, 1]$. We can write the price of anarchy as:

$$PoA = \frac{\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3}{\alpha_3 v_1 + \alpha_1 v_2 + \alpha_2 v_3}$$

Now, notice that the coefficient of v_2 is smaller in the numerator than in the denominator. The same is true for v_3 . Now, we use the following simple observation about ratios: if $a \leq b$ and $v \geq v'$ then: $\frac{x+av}{y+bv} \leq \frac{x+av'}{y+bv'}$, which is natural, because decreasing v we decrease the denominator more than the numerator. Now, we use that technique to bound v_2 and v_3 in terms of v_1 :

- $v_2 \geq \frac{\alpha_1 - \alpha_3}{\alpha_1} v_1$
- $v_3 \geq \frac{\alpha_2 - \alpha_3}{\alpha_2} v_1$

The first inequality comes from the Nash inequalities $\alpha_3(v_1 - 0) \geq \alpha_1(v_1 - b_2) \geq \alpha_1(v_1 - v_2)$ and the second comes from the fact that $\alpha_3(v_1 - 0) \geq \alpha_2(v_1 - b_3) \geq \alpha_2(v_1 - v_3)$. Now, we get:

$$PoA \leq \frac{\alpha_1 v_1 + \alpha_2 \left[\frac{\alpha_1 - \alpha_3}{\alpha_1} v_1 \right] + \alpha_3 \left[\frac{\alpha_2 - \alpha_3}{\alpha_2} v_1 \right]}{\alpha_3 v_1 + \alpha_1 \left[\frac{\alpha_1 - \alpha_3}{\alpha_1} v_1 \right] + \alpha_2 \left[\frac{\alpha_2 - \alpha_3}{\alpha_2} v_1 \right]} \quad (2)$$

Which allows us to eliminate v_1 and optimize for α . By standard techniques one can prove that the optimum is 1.25913 which is the root of a fourth degree equation. The values for which it is maximized are $\alpha_1 = 1, \alpha_2 = 0.55079, \alpha_3 = 0.4704$.

Case (ii): $\pi = [3, 1, 2]$. We can write the price of anarchy as:

$$PoA = \frac{\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3}{\alpha_2 v_1 + \alpha_3 v_2 + \alpha_1 v_3}$$

and again we use the same trick of realizing that $v_1 \leq \frac{\alpha_1}{\alpha_1 - \alpha_2} v_3$ by the fact that player 1 doesn't want to get the first slot, and $v_2 \leq \frac{\alpha_1}{\alpha_1 - \alpha_3} v_3$ by the fact that player 2 doesn't want to take the first slot. That gives us:

$$PoA \leq \frac{\alpha_1 \left[\frac{\alpha_1}{\alpha_1 - \alpha_2} v_3 \right] + \alpha_2 \left[\frac{\alpha_1}{\alpha_1 - \alpha_3} v_3 \right] + \alpha_3 v_3}{\alpha_2 \left[\frac{\alpha_1}{\alpha_1 - \alpha_2} v_3 \right] + \alpha_3 \left[\frac{\alpha_1}{\alpha_1 - \alpha_3} v_3 \right] + \alpha_1 v_3}$$

which has the same solution 1.25913 when maximized. Now, it is maximized for $\alpha_1 = 1, \alpha_2 = 0.5295, \alpha_3 = 0.1458$. In fact, it is not hard to see that those two PoA expressions have the same maximum: given a point $(1, \alpha_2, \alpha_3)$ (wlog we can consider $\alpha_1 = 1$ because the expression is homogeneous), the second expressions evaluates to the same value in the point $(1, 1 - \alpha_3, \frac{\alpha_2 - \alpha_3}{\alpha_2})$. ■

We proved that 1.25913 is the tight Price of Anarchy for 3 slots (we can use the optimization results in Case(i) to generate a tight example). We also conjecture that this is the correct Price of Anarchy for any $n \geq 3$. Moreover, we conjecture that the allocation maximizing the Price of Anarchy for n slots is $\pi = [2, 3, 4, \dots, n, 1]$, i.e., the player with higher value takes the bottom slot and all players $i > 1$ take slot $i - 1$. Then, if this is the case, we can prove our desired theorem by showing the following result:

Lemma 7 *If an equilibrium with n players and n slots is such that $\sigma(1) = n$ and $\sigma(i) = i - 1$ for the other players, then the Price of Anarchy is 1.25913.*

Proof. Following a proof scheme similar to used in Lemma 6 we can write:

$$PoA = \frac{\alpha_1 v_1 + \sum_{i>1} \alpha_i v_i}{\alpha_n v_1 + \sum_{i>1} \alpha_{i-1} v_i} \leq \frac{\alpha_1 v_1 + \sum_{i>1} \alpha_i \left[\frac{\alpha_{i-1} - \alpha_n}{\alpha_{i-1}} v_1 \right]}{\alpha_n v_1 + \sum_{i>1} \alpha_{i-1} \left[\frac{\alpha_{i-1} - \alpha_n}{\alpha_{i-1}} v_1 \right]}$$

This boils down to optimizing a function on multiple variables. It can be shown using standard techniques from optimization that the optimum is the same of equation 2. In fact, if $(\alpha_1, \alpha_2, 1)$ is a solution to 3 slots, then $(\alpha_1, \alpha_2, 1, \dots, 1)$ is a solution for n slots. ■

References

- [1] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *STOC08*, 2008.
- [2] Edelman, Benjamin, Ostrovsky, Michael, Schwarz, and Michael. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, March 2007.
- [3] Renato D. Gomes and Kane S. Sweeney. Bayes-nash equilibria of the generalized second price auction. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 107–108, New York, NY, USA, 2009. ACM.
- [4] J. Hannan. Approximation to bayes risk in repeated plays. In M. Dresher, A. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 4. Princeton University Press, 1957.
- [5] S. Kakade, A. Kalai, and K. Ligett. Playing games with approximation algorithms. In *STOC07*, 2007.
- [6] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 2005.
- [7] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.
- [8] Sébastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.
- [9] Renato Paes Leme and Éva Tardos. Pure and bayes-nash price of anarchy for generalized second price auctions. In *FOCS10: 51st Annual IEEE Symposium on Foundations of Computer Science*, 2010.

- [10] B. Lucier. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *ICS10: Innovations in Computer Science 2010*, 2010.
- [11] Brendan Lucier and Alan Borodin. Price of anarchy for greedy auctions. In *SODA '10*. ACM, 2010.
- [12] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.
- [13] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.
- [14] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522, New York, NY, USA, 2009. ACM.
- [15] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.